Using Pegasus™ for Advanced Node Physical Verification

Cadence Design Systems / Microsemi



ABSTRACT

Physical verification signoff challenges have exponentially grown from the 65nm node due to the escalating design size and complexity, complex design-rule checks (DRCs), higher DRC rule count and operations. At 28nm and 16nm, full-chip DRC runs time took multiple days on 64 CPUs. Adding more CPUs did marginally improve the turnaround time but beyond 128 there was little gain because the poor scalability of our incumbent tool.

Cadence introduced the Pegasus™ System to replace their previous Cadence Physical Verification System (PVS) product. With Pegasus, we have been able to scale full-chip DRC runs to hundreds of CPUs and achieve up to 10X speedup when compared to the incumbent tool. Pegasus gave us the flexibility to add hundreds of CPUs for peak usage and achieve our desired runtime.

Pegasus helps us regain the DRC productivity. We experienced near-linear scalability enabling more efficient utilization of our compute resources, jobs that previously took more than 40 hours to run have completed in less than three hours. In this paper, we will outline our evaluation results of Pegasus.



Using Pegasus™ for Advanced Node Physical Verification

Scott Barrick, Senior Manager, Advanced Engineering Services

Agenda

- Microsemi Overview
- Advanced Node Challenges
- Physical Verification System (PVS) at Microsemi
- Pegasus Evaluation
- Flow Migration
- Cloud Computing
- Summary

About Microsemi Corporation (Nasdaq: MSCC)

- Global provider of semiconductor solutions for applications focused on delivering power, reliability, security and performance.
- High-value, high barrier-entry markets:
 - Aerospace & Defense
 - Communications
 - Data Center
 - Industrial
- FY 2016 revenue: $1.65 billion



Corporate headquarters in Aliso Viejo, CA

Diversified, Defensible Markets

Aerospace & Defense *26% of Revenue*	Communications *36% of Revenue*	Data Center *23% of Revenue*	Industrial *15% of Revenue*
Military Communication	Wireless Backhaul	Hyperscale & Enterprise	Industrial Control
Information Assurance	Base Station	Servers	Medical Devices
Engine Control	Routing & Switching	Storage Systems	Energy
Avionics	Networking	Rack Disaggregation	Automotive
Electronic Actuation	Access & CPE	NVM Solutions	SEMI Cap Equipment
Bus & Payload Electronics	Wireless Terminal	Security	Machine to Machine
Launch Systems	Fiber Optic Backhaul		

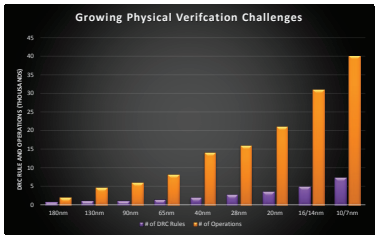
AIRBUS, LOCKHEED MARTIN, BOEING, CISCO, HUAWEI, ERICSSON, ZTE, SAMSUNG, NOKIA, Hewlett Packard Enterprise, DELL, IBM, Medtronic, GE, ST. JUDE MEDICAL

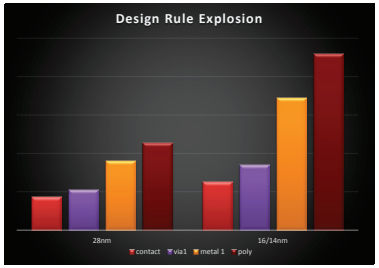
Based on FQ1 2017 revenue

Advanced Node Challenges

Microsemi
5

Design Rule Explosion

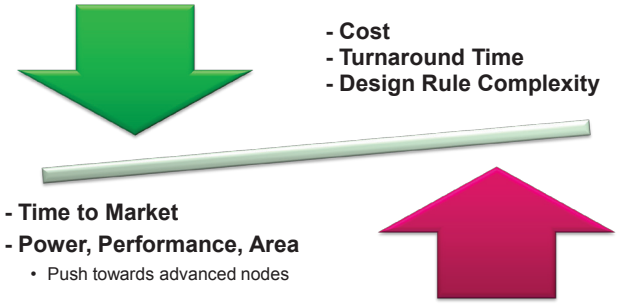




- Design rule explosion 65nm- to 40nm- to 28nm- to 16nm (over the years)
- New multi-order rules, now FinFET and coloring have added another dimension of complexity.
- Advanced node design requires multiple cycles to complete and requires enough cycle time to accommodate DRCs close to tapeout.

Microsemi
Power Matters.™

Advanced Node Constraints to Manage

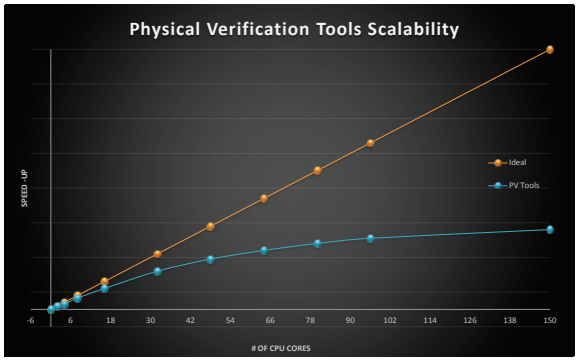


- Cost
- Turnaround Time
- Design Rule Complexity

- Time to Market
- Power, Performance, Area
- Push towards advanced nodes

Microsemi
Power Matters.™

Current DRC Tools are not Scalable



- Poor scalability of current tools require jobs to be scheduled in advance.
- Costly for last minute ECOs
- Need good scalability to be able to run DRC more often within design cycle, without raising costs

Microsemi
Power Matters.™

Physical Verification System (PVS) at Microsemi

Adopted PVS in 2012 for Physical Verification Signoff

- **Dozens of tape outs across TSMC 65nm, 40nm, 28nm**
 - Sizes ranging from 50mm^2 to 700mm^2
 - Variety of big D little A and big A little D devices
- **No rule deck issues**
 - Certified by foundries
- **No accuracy issues**
- **Actively designing on TSMC 16nm**
 - testchip tapeouts done.

Challenges with PVS

1. **Multi-CPU scalability diminishes after 64 cores**
 - This was fine in 65nm & 40nm.
 - For 28nm and below achieving overnight final DRC runs likely not be possible.
2. **Requires a "master" process to distribute across multiple servers.**
 - 4-8 CPUs that get things started, then sit idle for the bulk of the run
3. **Could not start until 100% of CPU resources were locked down.**
 - During tape-out crunch when the CPU farm is stressed you could secure 90% of CPUs and they would sit idle until the last CPUs were secured.

Pegasus™ Evaluation

Pegasus's Key Features

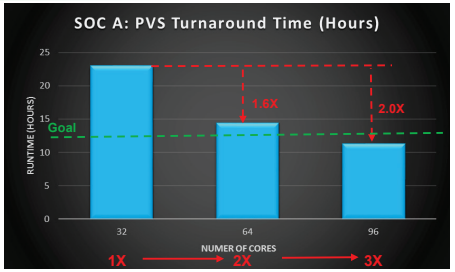
- Designed to achieve a high degree of scalability
 - Pipelined engines reduces the memory footprint
- Does not need a high memory "master" to launch jobs
 - Can use a small machines and launch many jobs for full chip runs
 - Ability to use disks that are "local" to the execution hosts
- Flexible distributed processing infrastructure allows easy acquisition of hosts
- Cloud Ready: Remote execution on internal or external clouds
- Built-in disk-swap prefers physical memory (faster) than writing to disk

Pegasus DRC Evaluation

- **Criteria:**
 - **Turnaround time:**
 - Minimum of overnight (< 12 hours), Ideally at least 5x faster than baseline
 - **Accuracy:** 100% match to PVS
 - **Efficient CPU utilization:** i.e double the CPUs, half the runtime
 - **Peak Usage:** Easy ramp up on burst of 100s of CPUs to get fast runtime
 - Without any need to dedicate or restrict hardware
- **Test cases:**

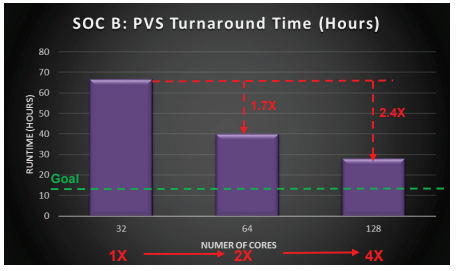
SOC A	SOC B
60M gates	400M gates
50M bits SRAM	400M bits SRAM
High Speed Serdes, DDR Phys	Many High Speed Serdes
GDSII: 9.8G gzipped	GDSII: 24G gzipped

Baseline PVS Runtime/Scalability



SOC A

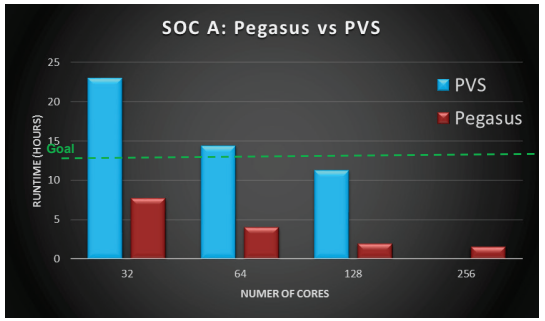
- Uses a 4-CPU master job to launch
 - Pure overhead
- Scaling 32 -> 64 CPU cut run by 37% OK
- Scaling 32 -> 96 CPU cut run by 51% OK
 - Did get us to our goal of overnight GOOD
- Scaling beyond 96 CPU ... Little benefit BAD



SOC B

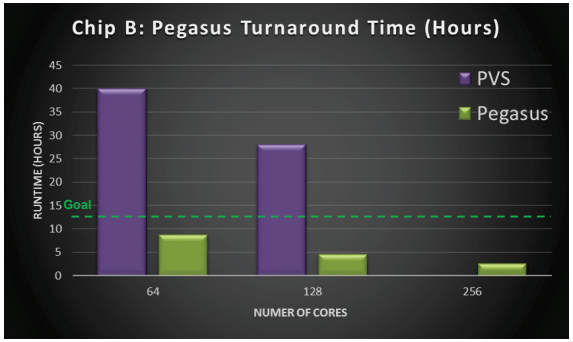
- Uses a 8-CPU master job to launch
 - Pure overhead
- Scaling 32 -> 64 CPU cut run by 40% OK
- Scaling 32 -> 128 CPU cut run by 58% OK
- Scaling beyond 128 CPU ... Little benefit BAD
- No amount of CPU could get us to overnight BAD

Runtime, Pegasus vs PVS: SOC A



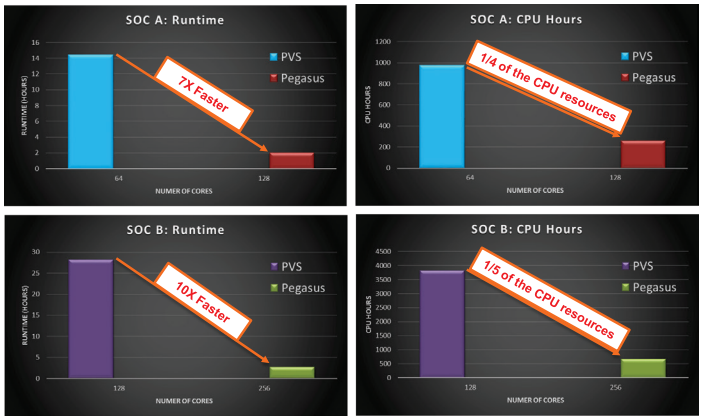
- **Pegasus Provides:**
 - 3.0-5.6x runtime improvement for same number of CPUs
 - >10x runtime improvement with increased CPUs
 - Linear scaling for 32-128 CPU
 - Scaling does flattens out on small designs using > 128 CPU

Runtime, Pegasus vs PVS: SOC B



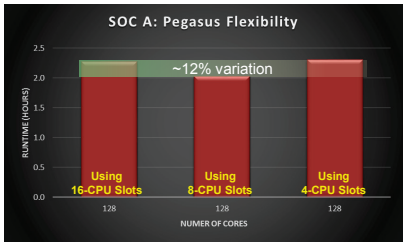
- **Pegasus Provides:**
 - 4.5-6x runtime improvement for same number of CPUs
 - > 10x runtime improvement with increased CPUs
 - Linear scaling for 64-256 CPU

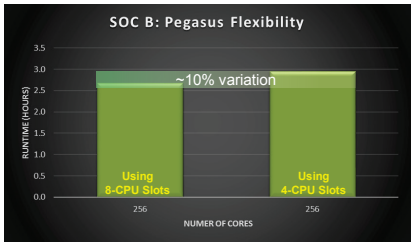
Efficiency, Pegasus vs PVS



- **Pegasus provides significant performance gains while reducing the load on the compute farm**

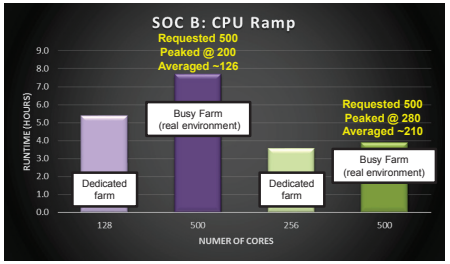
Pegasus Flexibility





- **Pegasus provides the flexibility to tailor jobs to better fit the current state of the compute farm without sacrificing performance**
 - As tape-out approaches it is typically much easier to find a large number of small CPU slots than a handful of larger CPU slots
- **Combine this with the near linear scaling allows easy adaptation of runs to hit the desired runtime and fit the current environment of the compute farm**

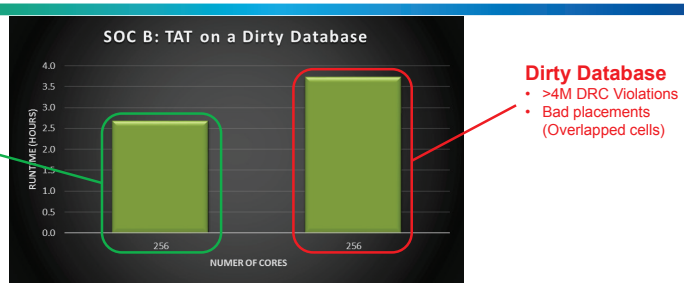
Pegasus CPU Ramp



- **The two "500 CPU" jobs were launched into a busy farm.**
 - This ensured they would not get 100% of the requested CPU
 - Had less control over the servers that the jobs landed and introduced more variability
- **Pegasus began as soon as the first CPUs were secured**
- **As LSF secures more CPUs Pegasus distributes work to them**

*Note : Pegasus version used here was older than that used for previous data so overall runtimes shown here are a little slower

Pegasus On A Dirty Database



- **The dirty database was the first fully assembled top level**
 - Overall turnaround time with Pegasus < 4hrs, which would allow multiple spins per day
 - Pegasus would have greatly accelerated the debug of the gross issues with the database
- **When this database was first run with PVS it had to be killed after 8 days.**
- **Pegasus performance gains seen even on extremely dirty database**

Flow Migration

Migrating from PVS to Pegasus



- **Decks worked, as is.**
- **The Pegasus configuration for LSF is different than PVS**
 - Required a few iterations to get it setup properly
- **That was it!**

Cloud Computing

High Performance Cloud Computing

- **Why Cloud computing is appealing?**
 - Internal solutions are slow to scale (weeks -> months):
 Quotes -> Purchase Req -> Approvals -> Purchase Order -> Shipping/Customs -> Install
 - Having sufficient hardware for large peaks means it is likely under utilized over the long term
 - Hardware purchased for peak #1 might be obsoleted by peak #2
- **Fast Access and Flexible (elasticity of resources) :**
 - Once initial setup is complete hardware can be added/removed very quick
 - Access to a wide variety of hardware so can better adapt on the fly to the design/application needs
 - Not stuck with the hardware purchased at the beginning of the project
 - Can attain much higher peak capacity

Cloud Computing at Microsemi

- **Multiple internal clouds located geographically to minimize latency to primary design sites**
 - Each sized to handle a large Rev A or multiple Rev Bs developments.
- **When needed external clouds are enabled in similar geographies as internal clouds**
 - Configured to have the same “look and feel” as internal cloud.
 - Jobs sent to a particular compute cluster (internal or external) have access to the same infrastructure regardless of location.
- **This model still requires a significant amount of planning/setup**
 - Network connectivity, local storage, encryption, data replication

But provides a low barrier of entry for users.

Pegasus Cloud Ready Computing

- **Massively parallel architecture showed near-linear scalability on hundreds of CPUs**
- **Plug and play setup allows users to run Pegasus easily on clouds**
 - Microsemi has evaluated on both internal and external clouds
- **Efficient CPU utilization across heterogeneous networks**
 - Combined with the scalability can easily adapt to various runtime and cost constraints.
 - With increasingly flexible cloud offerings could potentially turn usage to a “cost per turn” and allow design teams determine how many turns they are willing to pay for.

Summary

Pegasus DRC Evaluation

- Criteria:
 - ✓ **Turnaround time:**
 - Achieve 7-10X faster than baseline
 - ✓ **Accuracy:** 100% match to PVS
 - ✓ **Efficient CPU utilization:** i.e double the CPUs, half the runtime
 - Achieved near linear performance gains up 256 CPUs
 - ✓ **Peak Usage:** Easy ramp up on burst of 100s of CPUs to get fast runtime
 - Pegasus starts with first available CPUs and ramps as LSF secures more with no user intervention
- Test cases:

SOC A	SOC B
60M gates	400M gates
50M h…	400M … M
…peed Serdes, DDR Phys	…igh Speed Serdes
GDSII: 9.8G gzipped	GDSII: 24G gzipped

Can be run in <2hrs

Can be run in <3hrs

Challenges with PVS

1. **Multi-CPU scalability diminishes after 64 cores**

Addressed in Pegasus

2. **Requires a "master" process to distribute across multiple servers.**

Eliminated in Pegasus

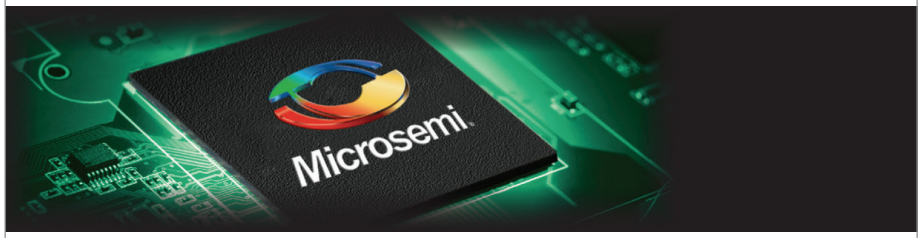
3. **Could not start until 100% of CPU resources were locked down.**

Eliminated in Pegasus

Pegasus Summary

- **All Microsemi evaluation criteria were met**
- **All Microsemi challenges with PVS are addressed or eliminated**
- **DRC moves from a painful necessity to an easy final step**
 - Demonstrated impressive scalability and dramatic reduction in turnaround time
 - Pegasus ushers in a new DRC use model of running blocks over coffee and lunch breaks, and full chip faster than overnight.
- **Cloud-ready processing**
 - Futuristic technology which will drive more conversation between IC companies and the foundries
 - Promising to scale up resources at short notice during peak tapeout
- **100% accurate and uses foundry PVS rule decks**

Power Matters



Thank you